

Article

Not peer-reviewed version

A Sustainability-Aware Federated Graph Attention Framework for Supply Chain Process Modelling

[Vasileios Alexiadis](#), [Maria Drakaki](#)^{*}, [Panagiotis Tzionas](#)

Posted Date: 28 January 2026

doi: 10.20944/preprints202601.2211.v1

Keywords: sustainability-aware federated learning; graph attention networks; federated graph neural networks; supply chain process modelling; privacy-preserving analytics; carbon-aware message passing



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Sustainability-Aware Federated Graph Attention Framework for Supply Chain Process Modelling

Vasileios Alexiadis ¹, Maria Drakaki ^{1,*} and Panagiotis Tzionas ²

- ¹ Department of Science and Technology, University Center of International Programmes of Studies, International Hellenic University, 14th Km Thessaloniki-N. Moudania, GR-57001 Thermi, Greece
- ² Department of Industrial Engineering and Management, International Hellenic University, P.O. Box 141, GR-57400 Thessaloniki, Greece
- * Correspondence: mdrakaki@ihu.edu.gr; Tel.: +30-2310-807524

Abstract

Modern supply chains operate as highly interconnected networks characterized by decentralization, data silos, and increasing sustainability constraints. While graph neural networks (GNNs) have demonstrated strong performance in modelling such relational systems, their practical deployment is hindered by limited data sharing across organizations. Federated learning (FL) offers a promising solution by enabling collaborative model training without exposing proprietary data, yet existing approaches rarely integrate graph structure and sustainability objectives simultaneously. This study proposes a sustainability-aware federated graph attention framework for supply chain process modelling. The framework combines Graph Attention Networks (GATs) with federated optimization to learn from decentralized, partially observable supply chain subgraphs while embedding environmental considerations directly into the attention mechanism. A synthetic multi-tier supply chain case study is developed to evaluate the approach under realistic data-governance constraints. Experimental results show that while centralized graph learning achieves the highest predictive accuracy, the proposed sustainability-aware federated GAT attains competitive performance relative to standard federated baselines, while systematically reducing reliance on carbon-intensive transport links. An ablation analysis demonstrates a smooth and controllable trade-off between predictive accuracy and sustainability alignment through a single policy parameter. The findings highlight the feasibility of privacy-preserving, sustainability-informed graph learning for supply chain process modelling and provide a principled foundation for environmentally aligned AI deployment in multi-enterprise settings.

Keywords: sustainability-aware federated learning; graph attention networks; federated graph neural networks; supply chain process modelling; privacy-preserving analytics; carbon-aware message passing

1. Introduction

Foundational work in industrial systems modelling established that production, warehousing, and supply-chain operations are dynamic, interdependent, and event-driven, motivating formal modelling and intelligent control. Agent-based systems combined with Colored Petri Nets (CPNs) were shown to capture this complexity effectively, improving coordination in warehouse order picking and resource allocation under demand uncertainty [1]. These approaches were extended to multi-echelon supply chains, where hierarchical CPNs enabled analysis of inventory propagation, order variability, and disruption effects across tiers [2,5]. Empirical studies further demonstrated that inventory inaccuracy significantly degrades service levels and amplifies upstream instability, highlighting the need for system-wide coordination [3]. Reinforcement learning integrated with timed CPNs enabled adaptive manufacturing scheduling under stochastic disturbances, reinforcing the view of supply chains as networked decision systems with cascading effects [4].

In parallel, intelligent multi-agent systems were applied to production maintenance and asset health monitoring. Distributed architectures combining machine-learning classifiers and autonomous agents achieved high fault-detection accuracy in industrial equipment such as induction motors [6,7], demonstrating that decentralized intelligence can support predictive maintenance while preserving modularity and data locality. Subsequent conceptual work framed production and maintenance as a unified learning ecosystem, emphasizing coordination, shared representations, and adaptive decision-making across operational functions [8]. Together, these studies laid the groundwork for decentralized, learning-driven industrial systems operating over structured interaction networks.

Building on this trajectory, modern supply chains are increasingly modeled as large-scale networks of suppliers, manufacturers, distributors, and retailers, where efficiency, resilience, and sustainability are critical. AI-based process modelling has therefore become central to supply-chain optimization. In this context, Graph Neural Networks (GNNs) are particularly suitable, as supply chains can be naturally represented as graphs with entities as nodes and material, informational, or financial dependencies as edges [9]. Empirical evidence shows that GNN-based models outperform conventional statistical and deep-learning baselines on tasks such as forecasting and classification, often achieving 10–30% improvements on key metrics by learning non-linear structural dependencies and propagation effects [9].

Attention-based graph models further enhance supply-chain learning. Graph Attention Networks (GATs) assign non-uniform weights to neighboring nodes during aggregation, enabling models to emphasize critical supplier–customer relationships while suppressing less relevant interactions [10]. Beyond predictive gains, attention mechanisms improve interpretability by revealing influential dependencies, and recent GAT-based resilience models demonstrate strong disruption-classification performance while exposing structural vulnerabilities through learned attention patterns [11].

Despite these advances, real-world deployment is constrained by data decentralization and confidentiality. Supply-chain data are distributed across organizations and are rarely shareable due to commercial, contractual, and regulatory restrictions, rendering centralized GNN training impractical in multi-enterprise settings [12]. Federated Learning (FL) addresses these constraints by enabling collaborative model training without exchanging raw data. The foundational Federated Averaging approach demonstrated that high-quality neural models can be trained across decentralized, non-IID data while retaining data locally [13]. In industrial ecosystems, FL enables data-limited manufacturers and logistics partners to benefit from collective learning under confidentiality constraints [14].

A further driver shaping supply-chain analytics is the need to integrate sustainability objectives. Supply chains contribute significantly to greenhouse gas emissions and resource use, motivating sustainability-aware AI that incorporates metrics such as carbon emissions or energy intensity directly into learning and decision-making. AI-driven logistics and operational optimization have been associated with measurable environmental benefits, including substantial CO₂ reductions [15]. In adjacent domains, federated learning combined with control and optimization—such as reinforcement learning in energy systems—has achieved simultaneous cost and emissions reductions through decentralized coordination [16]. Nonetheless, most supply-chain ML systems remain predominantly single-objective, and explicit sustainability integration during training is still uncommon, particularly under decentralization constraints.

Consequently, important gaps remain. Centralized GNNs struggle with data silos and governance, while conventional FL is typically designed for tabular data and does not naturally accommodate graph-structured dependencies spanning organizational boundaries. Only recently has federated graph learning been applied to supply-chain visibility, risk modelling, and demand–inventory prediction, including early attempts to integrate sustainability considerations [17,18]. These studies demonstrate feasibility but also underscore the absence of a unified framework jointly

addressing decentralization, graph-based interdependence modelling, and sustainability-aware optimization.

This paper addresses this gap by proposing a sustainability-aware federated graph attention framework for supply-chain process modelling. We develop a federated GAT architecture suitable for multi-enterprise supply chains, integrate sustainability signals directly into the learning objective, and evaluate the approach on a realistic multi-tier case study against centralized and non-collaborative baselines, analyzing both predictive performance and sustainability-aligned behavior under strict data-governance constraints.

The remainder of the article is organized as follows. Section 2 reviews related work on GNNs, federated learning, sustainability-aware AI, and federated graph learning. Section 3 presents the methodology and system architecture. Section 4 reports experimental results. Section 5 discusses implications and limitations, and Section 6 concludes with directions for future research.

2. Related Work

To contextualize our proposed framework, we survey relevant literature in four key areas. First, we discuss how graph neural networks have been applied in supply chain modelling and optimization, highlighting their benefits over traditional methods. Next, we review federated learning approaches in industrial domains, focusing on how privacy-preserving collaboration has been enabled in manufacturing and supply chain scenarios. Third, we examine research on sustainability-aware AI methods that integrate environmental objectives into model training or decision-making, including multi-objective optimization techniques. Finally, we cover the relatively few works that have attempted to combine GNNs with federated learning, noting their achievements and limitations.

2.1. Graph Neural Networks in Supply Chain Management

Supply chains are inherently graph-structured systems, with nodes representing firms or facilities and edges encoding material, contractual, transportation, or information dependencies. This structure makes Graph Neural Networks (GNNs) particularly suitable, as they learn over non-Euclidean data by propagating information across network topology. Unlike independent-sample models, GNNs capture higher-order interdependencies and cascading effects, enabling them to model how local disruptions propagate across multiple supply chain tiers through message passing [9].

Empirical evidence consistently shows that GNNs outperform classical time-series and conventional machine-learning approaches in supply chain tasks such as demand forecasting, inventory control, risk assessment, anomaly detection, and network reconstruction. Wasi et al. report performance gains typically in the range of 10–30%, attributable to GNNs' ability to exploit shared structural context (e.g., common upstream suppliers) and enable cross-entity generalization [9]. Industrial studies further confirm these advantages, particularly under disruption scenarios where structural dependencies dominate system behavior, as demonstrated in automotive supply networks by Gupta et al. [19].

Beyond accuracy, GNNs provide structural interpretability, which is increasingly critical for decision support, risk governance, and regulatory compliance. Graph Attention Networks (GATs) enhance this capability by learning adaptive attention weights over neighbors, allowing models to emphasize influential entities and relationships. Attention-based graph models have been shown to improve both predictive performance and interpretability in large-scale supply networks, disruption prediction, and manufacturing resilience applications [20–22].

GNNs are also effective for uncovering latent or incomplete supply chain structure. Prior work formulates supply chain visibility as a link-prediction problem, showing that GNNs can infer hidden supplier relationships and indirect dependencies absent from enterprise records [23]. Extensions that integrate graph embeddings with heterogeneous data sources, such as textual risk signals, further improve systemic risk identification in ICT and critical infrastructure supply chains [24].

Despite these advances, most GNN-based supply chain studies assume access to a centralized and fully observable supply chain graph. In practice, this assumption is often invalid due to confidentiality, data-protection regulation (e.g., GDPR), contractual restrictions, and competition law, which collectively limit cross-enterprise data sharing. Additionally, existing GNN applications predominantly optimize predictive or operational performance, while sustainability objectives—such as emissions intensity or energy use—are typically treated as ex-post evaluation criteria rather than embedded learning targets.

These limitations motivate decentralized graph-learning approaches that operate under partial visibility and organizational boundaries, while supporting multi-objective optimization aligned with emerging sustainability and regulatory requirements. This gap is addressed in subsequent sections through the integration of federated learning, graph attention mechanisms, and sustainability-aware modelling objectives.

2.2. Federated Learning in Industrial and Supply Chain Settings

Federated Learning (FL), introduced by McMahan et al. through Federated Averaging, enables effective training of deep neural networks across decentralized and non-IID data without sharing raw data [13]. Subsequent surveys identify FL as particularly suitable for cross-silo industrial environments characterized by strong heterogeneity in scale, data quality, and processes [25]. In Industrial IoT settings, hierarchical and two-stage aggregation strategies improve convergence stability and reduce communication overhead, supporting factory-scale deployment [26].

Within Industry 4.0 and Industry 5.0 frameworks, FL is increasingly regarded as a core enabler of privacy-preserving collaborative intelligence. Surveys in smart manufacturing and product lifecycle management demonstrate its applicability to quality control, predictive maintenance, and process optimization while maintaining data locality [27]. Empirical studies further show that FL enables data-limited manufacturers to achieve performance comparable to centralized learning without exposing proprietary data, as demonstrated in additive manufacturing condition monitoring [28]. These properties align with emerging trustworthy-AI principles emphasizing privacy, auditability, and human-centric governance [29].

In supply chain management, FL applications remain limited but are expanding. Nguyen et al. propose a federated framework for delivery-delay prediction across textile suppliers, demonstrating feasible privacy-preserving collaboration and improved generalization relative to single-organization baselines [30]. The same study shows that federated models can rival or exceed centralized performance by mitigating overfitting across heterogeneous participants [30]. Related work extends FL to cross-border logistics, enabling early-warning systems for disruption risk while respecting jurisdiction-specific data-sovereignty constraints [31]. Comparable federated demand-forecasting architectures in retail and agri-food supply chains rely on encrypted model updates and align with emerging sectoral data-governance frameworks [32,33].

Despite its promise, cross-silo FL introduces substantial system and governance challenges, including coordination across autonomous organizations, heterogeneous compute resources and data schemas, and exposure to inference attacks. Secure aggregation and cryptographic mechanisms are therefore required in adversarial or semi-honest settings [25]. Tang et al. address these challenges by integrating FL with graph neural networks for privacy-preserving supply-chain data sharing [34]. At the organizational level, FL deployments must also comply with data-protection, trade-secret, and competition regulations, which may restrict even indirect information leakage.

Overall, FL is increasingly recognized as a cornerstone technology for collaborative industrial AI, yet supply-chain adoption remains largely exploratory. Existing applications predominantly rely on conventional neural architectures and tabular representations, despite the inherently relational structure of supply chains. This gap motivates the integration of FL with graph-based and attention-driven models to enable privacy-preserving, structure-aware learning across multi-enterprise supply networks.

2.3. Sustainability-Aware Artificial Intelligence

As sustainability targets intensify, AI research distinguishes between AI-for-sustainability, which seeks to reduce emissions, energy use, and resource intensity in operational systems, and “green” AI, which focuses on minimizing the environmental footprint of AI itself through efficient training, inference, and reporting. This work primarily addresses the former by embedding environmental objectives directly into AI models for supply chain process management, while accounting for computational and communication costs.

In supply chain management, sustainability has become a central operational concern, as supply-chain activities contribute significantly to greenhouse gas emissions and resource use, requiring the joint optimization of environmental, cost, and service objectives [35]. Accordingly, indicators such as carbon footprint, energy intensity, waste, and circularity are increasingly treated as decision variables rather than ex-post reporting metrics. AI methods are particularly suited to this setting, as high-dimensional operational data and complex constraints limit analytical optimization. Empirical evidence shows that AI-driven analytics can improve green supply-chain process integration and environmental performance when aligned with operational governance [35], while systematic reviews identify logistics optimization, demand–inventory coordination, and risk-aware planning as the most mature application areas, where emissions and energy use are tightly coupled to operational decisions [36].

A core methodological principle in sustainability-aware AI is treating environmental performance as a first-class learning objective. This includes integrating environmental impacts into loss functions, enforcing emissions constraints, or learning policies that explicitly trade off operational and environmental KPIs under uncertainty. In supply chains, Abushaega et al. formalize this approach through a multi-objective framework combining federated learning and graph neural networks to jointly optimize operational and environmental objectives [18]. Related work in energy systems demonstrates that decentralized learning can yield simultaneous cost and emissions reductions, including empirically observed savings under federated reinforcement learning for building energy management [16]. Additional studies show that federated sequence models support sustainability optimization beyond centralized deployments [37,38].

Sustainability-aware AI also encompasses measurement and accountability of AI’s own environmental footprint. The literature emphasizes standardized green evaluation metrics for comparability across models and deployments. Borraccia et al. propose hybrid metrics for assessing energy–carbon impacts of AI pipelines [39], while bibliometric analyses document rapid growth in sustainability-oriented machine-learning research alongside a persistent gap between methodological advances and deployment-level impact [40,41]. Organizational studies further indicate that effective AI-for-sustainability outcomes depend on governance structures and integration into decision workflows, not solely on model design [42].

Finally, sustainability benefits from AI in supply chains are not guaranteed and may be offset by rebound effects, misaligned incentives, or increased computational burden. Practitioner reports often claim large emissions reductions from AI-driven logistics and planning interventions [43], but the research literature stresses rigorous evaluation, transparent baselines, and robust measurement to distinguish genuine environmental improvements from confounded operational effects [39,44]. While the feasibility of sustainability-aware AI is well supported, most supply-chain ML systems still prioritize operational accuracy, treating sustainability as an external constraint. The novelty of our approach lies in explicitly encoding sustainability signals within a decentralized federated graph-learning framework, enabling balanced optimization of operational and environmental KPIs under realistic data-governance constraints.

2.4. Combining Graph Neural Networks and Federated Learning

Bringing together GNNs and FL poses both conceptual and system-level challenges, and only a limited body of work has explored this intersection. Standard FL formulations implicitly assume that data points are independent and identically distributed across clients. Graph-structured data violates

these assumptions in two fundamental ways: (i) observations (nodes and edges) are relationally coupled, and (ii) the graph topology itself may be partitioned across organizational boundaries. In federated settings, this implies that informative dependencies can span clients, while each participant only observes a local subgraph. These properties complicate gradient aggregation, convergence, and representation learning. Despite these difficulties, early studies demonstrate that federated optimization can be extended to graph-based models, giving rise to the paradigm of Federated Graph Neural Networks (FedGNN).

The earliest line of work has been driven by privacy-sensitive applications such as recommender systems and financial networks. Wu et al. introduced FedGNN for privacy-preserving recommendation, training GNNs over decentralized user–item interaction graphs while applying local privacy mechanisms to gradients and structural signals [38]. Their results show that a federated GNN can achieve recommendation quality comparable to centralized baselines, while significantly reducing information leakage under realistic threat models. This finding is important because it establishes that graph-based representation learning remains viable under strict data-locality constraints, even when the global graph is never materialized.

Subsequent work extends federated graph learning to spatio-temporal domains. Meng et al. propose a cross-node federated GNN for traffic forecasting, where each client manages a regional subgraph and the global model captures dependencies that span geographic boundaries [38]. Their framework improves forecasting accuracy relative to purely local models and introduces mechanisms for mitigating non-IID graph distributions. Complementary research addresses *vertical* graph partitioning, in which different clients hold disjoint feature sets or edge views for the same nodes. Chen et al. propose a vertically federated GNN (VFGNN) that enables privacy-preserving node classification when features, edges, and labels are distributed across parties [39]. These works collectively demonstrate that both horizontal and vertical graph partitioning can be accommodated within federated optimization, albeit with additional architectural and cryptographic complexity.

Our framework distinguishes itself by introducing a Graph Attention Network within a federated learning architecture tailored to multi-enterprise supply chains. Attention is particularly valuable in federated graph settings because local subgraphs differ in topology, density, and informational relevance. An attention-based aggregator can dynamically weight inter-node and inter-client dependencies, allowing the global model to adapt to heterogeneous structural contexts more flexibly than fixed-weight GCNs. Moreover, we explicitly integrate sustainability objectives into the federated GNN training process—an element largely absent from existing FedGNN literature, which typically prioritizes predictive performance and privacy guarantees alone. While Abushaega et al. take an important step toward multi-objective federated graph learning for supply chains [11], their architecture does not exploit attention mechanisms and does not explicitly address the interpretability and cross-client structural asymmetries inherent in real-world networks.

3. Methodology and Experimental Setup

This section describes the experimental design used to evaluate the proposed federated sustainability-aware graph attention framework. The setup is deliberately constructed to reflect key structural and data-governance characteristics of real industrial supply chains, while remaining fully reproducible and analytically controlled.

The proposed federated GAT model is trained using an iterative federated learning process, which is designed to respect the data locality of each client while yielding a robust global model. The training follows the standard Federated Averaging (FedAvg) workflow, composed of repeated communication rounds between the central server and the distributed clients. Algorithm 1 summarizes the training procedure in pseudocode form.

Algorithm 1. Federated Training Workflow

Input: Initial global model parameters W^0 ; total rounds T ; clients $1, \dots, K$ each with local graph data.

Output: Trained global model parameters W^T .

- Initialize: Server initializes model weights W^0 and sends W^0 to all client nodes.
- For each round $t = 1$ to T :
 - a. *Client-Side Training*: Each client k receives current global weights W^{t-1} . Client k trains the GAT model on its local subgraph data (for several epochs or until local convergence) using W^{t-1} as initialization. Let ΔW_k^t be the weight updates or the new weights obtained on client k .
 - b. *Model Update Sharing*: Each client k sends its updated model (or just the weight difference) ΔW_k^t to the server. (No raw node/edge data or local labels are shared – only model parameters).
 - c. *Server-Side Aggregation*: The server aggregates the received updates from all clients to form a new global model. Typically, this is done by weighted averaging: $W^t \leftarrow \sum_{k=1}^K \frac{n_k}{N} (W^{t-1} + \Delta W_k^t)$, where n_k is the number of training samples (or another proxy for data size) at client k , and $N = \sum_k n_k$. This is the FedAvg update rule [5]. The server then distributes the updated global weights W^t back to all clients for the next round.
- End For (after T rounds or when a convergence criterion is met).
- Return the final global model W^T .

3.1. Synthetic Supply Chain Simulation

We evaluate the proposed federated GAT on a realistic synthetic multi-tier supply chain case study. The supply network (illustrated schematically in Figure 1) is modeled as a directed graph $G = (V, E)$ with multiple echelons including suppliers, manufacturers, distribution centers, and retailers. Each node represents an entity (e.g., a factory, warehouse, or retail outlet) and each directed edge represents a supply relationship or material flow between entities. Nodes are assigned to tiers by operational role (upstream suppliers in Tier 1, intermediate manufacturers in Tier 2, distribution centers in Tier 3, and retailers in Tier 4), creating a layered structure analogous to a consumer goods supply chain.




Figure 1. Multi-tier supply chain graph with operational (size/width) and sustainability (color/alpha) attributes.

In the implementation, tier sizes are explicitly defined 10 suppliers (S1–S10), 6 manufacturers (M1–M6), 4 distribution centers (D1–D4), and 8 retailers (R1–R8), for a total of $N = 28$ nodes. Directed edges are generated only between adjacent tiers ($S \rightarrow M$, $M \rightarrow D$, $D \rightarrow R$), producing a multi-tier dependency structure that naturally yields converging flows (multiple suppliers feeding a manufacturer) and diverging flows (a distributor feeding multiple retailers). Concretely, the stochastic edge generator draws for each source node an out-degree:

$$k_s = \max(1, \lfloor \mathcal{N}(\mu_k, 1) \rfloor), \quad (1)$$

then connects to a uniformly sampled subset of downstream nodes. This construction yields sparse, tier-constrained connectivity consistent with hierarchical supply networks.

To ensure that downstream tiers are not disconnected from upstream supply, we enforce a weak connectivity constraint, where for every node v in tiers 2–4, at least one inbound edge must exist. Formally, if

$$\deg^-(v) = 0, \quad (2)$$

a random inbound edge (u, v) is added from the upstream tier. This prevents isolated nodes and ensures that message passing is well-defined for most nodes in the node-regression task.

Each node $v \in V$ is assigned operational and sustainability proxy attributes drawn from tier-specific Gaussian distributions, namely capacity, CO₂ intensity, and energy intensity. These distributions are explicitly encoded inducing systematic heterogeneity across tiers (e.g., manufacturers have higher mean carbon intensity than retailers) while maintaining stochastic variability under a fixed random seed SEED = 42. Each edge $(u, v) \in E$ is parameterized by flow and distance (also positive truncated normal draws), and a deterministic transport-emissions proxy:

$$\text{transport-emissions}_{uv} = \text{flow}_{uv} \cdot \text{distance}_{uv} \cdot 10^{-3}. \quad (3)$$

This proxy is used consistently both for feature construction and for sustainability-aware attention modulation.

To simulate federated data partitioning, we divided the global supply chain graph into subgraphs held by distinct client organizations. Figure 2 illustrates this partitioning: each company (or region) in the supply chain acts as a federated client that observes only its portion of the graph. In the implementation, 4 federated clients are created.

Client assignment is performed independently within each tier using a seeded shuffle followed by round-robin allocation:

$$\text{client}(v_i) = (i \bmod K) + 1. \quad (4)$$

This produces tier-balanced client memberships but non-identical local structures. Edges that cut across clients represent inter-company links (e.g., a supplier shipping to a manufacturer in a different organization) – these links exist in the overall graph but no single client can exploit them during local training. This setup reflects real-world data silos: each enterprise has visibility into its local operations and direct partners, but not the full end-to-end supply chain.

Critically, the federated training loop enforces strict locality: each client trains on its induced subgraph consisting only of nodes it owns and intra-client edges. Formally, client k observes:

$$G_k = (V_k, E_k), E_k = \{(u, v) \in E : u \in V_k \wedge v \in V_k\}. \quad (5)$$

This is implemented by filtering the global edge list. Hence, cross-client dependencies exist globally but are missing locally, creating a non-IID, partial-observability regime.




Figure 2. Federated client–server representation of the supply chain graph.

3.2. Data Generation and Sustainability Labeling

We generate synthetic operational data on the supply chain graph to train and evaluate the models. Each node is assigned a feature vector representing its operational state and sustainability

profile, and each edge carries logistics attributes including transport-emissions. In the code, the node-level feature vector is explicitly defined as a 5-dimensional descriptor:

$$\mathbf{x}_v = [\text{capacity}_v, \text{co2_intensity}_v, \text{energy_intensity}_v, \text{total_in_flow}_v, \text{total_in_emis}_v] \in \mathbb{R}^5, \quad (6)$$

where the last two components summarize upstream dependency and upstream logistics footprint:

$$\text{total_in_flow}_v = \sum_{u:(u,v) \in E} \text{flow}_{uv}, \text{total_in_emis}_v = \sum_{u:(u,v) \in E} \text{transport_emissions}_{uv}. \quad (7)$$

This construction makes the prediction problem structurally dependent: the feature vector is not purely local, but contains aggregated information from incoming edges, which graph-based learners can exploit more naturally than independent baselines.

The feature matrix $X \in \mathbb{R}^{N \times 5}$ is standardized using dataset-level z-scoring:

$$X_{\text{norm}} = \frac{X - \mu}{\sigma + 10^{-9}} \quad (8)$$

and converted into the tensor X_t for training on CPU/GPU depending on availability.

Using this data, we formulate a synthetic target variable at each node called the “stress score.” This node-level stress score is designed as a composite outcome that increases under capacity scarcity, upstream dependency, and carbon-intensive logistics. In the implementation, the pre-normalized stress score $\tilde{y}_v$ is constructed as a weighted combination of interpretable drivers:

$$\tilde{y}_v = 0.45 \cdot \frac{1}{\text{capacity}_v + \varepsilon} + 0.20 \cdot \frac{\text{total_in_flow}_v}{\text{total_in_flow} + \varepsilon} + 0.25 \cdot \frac{\text{total_in_emis}_v}{\text{total_in_emis} + \varepsilon} + 0.10 \cdot \frac{\text{co2_intensity}_v}{\text{co2_intensity} + \varepsilon} + 0.05 \cdot \frac{\text{energy_intensity}_v}{\text{energy_intensity} + \varepsilon}, \quad (9)$$

with ε used only for numerical stability. The value is min-max scaled to $[0, 1]$ and perturbed with Gaussian noise of standard deviation $\text{std} = 0.05$:

$$y_v = \text{clip}(\text{minmax}(\tilde{y}_v) + \eta, 0, 1), \eta \sim \mathcal{N}(0, \sigma^2). \quad (10)$$

By construction, this target is influenced by both operational dynamics (capacity, inbound dependency) and environmental factors (carbon/energy intensity and transport emissions), aligning with the objective of sustainability-aware modelling.

We split the dataset into training, validation, and test sets by randomly partitioning the node set (not multiple temporal snapshots). Specifically, a random permutation of node indices is drawn and masks are created using train ratio = 0.7 and validation ratio = 0.15, test is the remainder. This yields a transductive node regression setting where the graph structure and all nodes are present during training, but only training nodes contribute to the supervised loss, while validation and test nodes are used strictly for evaluation.

3.3. Model Configurations and Training Procedure

We compare four model configurations in our experiments, corresponding to baseline approaches and our proposed method.

1. *Centralized GAT*: A Graph Attention Network model that has access to the entire supply chain graph during training. It represents an ideal pooled-data scenario and provides an upper-bound reference under full observability. The implemented GAT uses a single attention layer with 32 hidden dimensions and 2 heads. Message passing is performed over incoming neighborhoods. For head h , node projections are computed as:

$$\mathbf{h}_v^{(h)} = W^{(h)} \mathbf{x}_v, \quad (11)$$

and attention logits for an incoming edge (u, v) are computed using concatenated projected states and edge features $\mathbf{e}_{uv} \in \mathbb{R}^2$ (flow and transport emissions, min-max normalized over global edges):

$$\mathbf{e}_{uv}^{(h)} = \text{LeakyReLU}(\langle \mathbf{a}^{(h)}, [\mathbf{h}_v^{(h)} \parallel \mathbf{h}_u^{(h)} \parallel \mathbf{e}_{uv}] \rangle). \quad (12)$$

Normalization is performed by softmax over $u \in \mathcal{N}_{\text{in}}(v)$:

$$\alpha_{uv}^{(h)} = \frac{\exp(\mathbf{e}_{uv}^{(h)})}{\sum_{u' \in \mathcal{N}_{\text{in}}(v)} \exp(\mathbf{e}_{u'v}^{(h)})}. \quad (13)$$

The head outputs are aggregated by weighted neighbor sums and concatenated before a linear readout.

2. *Federated GAT (FedAvg)*: A standard federated version of the GAT model trained with synchronous Federated Averaging. Each client trains on its local induced subgraph $G_k = (V_k, E_k)$ for 10 local epochs over 40 rounds. After each round, client models are aggregated using node-count weights:

$$\theta^{(t+1)} = \sum_{k=1}^K \frac{|V_k|}{\sum_{j=1}^K |V_j|} \theta_k^{(t)}. \quad (14)$$

This baseline isolates the impact of partial observability and non-IID graph partitioning without additional sustainability mechanisms.

3. *Federated Sustainability-Aware GAT*: Our proposed model, which extends the federated GAT with a sustainability-aware attention mechanism. In the implementation, sustainability is incorporated directly into the attention coefficients. After computing $\alpha_{uv}^{(h)}$ by softmax, the model applies an exponential penalty based on normalized transport emissions $\text{emis_norm}_{uv} \in [0, 1]$ (the second edge-feature component):

$$\alpha_{uv}^{(h)} \leftarrow \alpha_{uv}^{(h)} \cdot \exp(-\lambda \cdot \text{emis_norm}_{uv}), \quad (15)$$

followed by re-normalization over the incoming neighborhood:

$$\alpha_{uv}^{(h)} \leftarrow \frac{\alpha_{uv}^{(h)}}{\sum_{u' \in \mathcal{N}_{\text{in}}(v)} \alpha_{u'v}^{(h)} + 10^{-9}}. \quad (16)$$

The supervised objective remains mean squared error on node stress prediction, i.e.,

$$\mathcal{L}_k = \frac{1}{|T_k|} \sum_{v \in T_k} (\hat{y}_v - y_v)^2, \quad (17)$$

where $T_k \subseteq V_k$ are the client's training nodes defined by the global train mask restricted to V_k .

All models were implemented in Python 3.11.9 using PyTorch and trained with the Adam optimizer using learning rate $\text{LR}=10^{-3}$ and weight decay $\text{WD}=10^{-4}$. Centralized models are trained for 250 epochs. Federated models are trained for 40 rounds with 10 local epochs per round, yielding 400 local epochs per client in total, while global validation RMSE is monitored each round by evaluating the aggregated model on the full graph and the global validation mask.

3.4. Evaluation Metrics

We evaluate model performance on the held-out test nodes using standard regression metrics: Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) between the predicted and true stress scores. For a test set $\mathcal{T} \subseteq V$, these are:

$$\text{RMSE} = \sqrt{\frac{1}{|\mathcal{T}|} \sum_{v \in \mathcal{T}} (\hat{y}_v - y_v)^2}. \quad (18)$$

$$\text{MAE} = \frac{1}{|\mathcal{T}|} \sum_{v \in \mathcal{T}} |\hat{y}_v - y_v|. \quad (19)$$

In addition to prediction error, a key evaluation in our study is the attention–emission relationship. Since the GAT explicitly outputs per-edge attention values, we can quantify sustainability alignment by pairing each learned attention coefficient with its corresponding normalized edge emission. Concretely, with attention extracted as head-averaged values.

$$\bar{\alpha}_{uv} = \frac{1}{H} \sum_{h=1}^H \alpha_{uv}^{(h)}, \quad (20)$$

and emissions given by emis_norm_{uv} , we compute rank-based associations (e.g., Spearman correlation) between $\bar{\alpha}_{uv}$ and emis_norm_{uv} , and complementary diagnostics such as the fraction of attention mass assigned to the highest-emission quantile of incoming edges. These measures are defined at setup level here, while empirical values and comparative analysis are reported in the results section.

4. Results

4.1. Predictive Performance Comparison

Table 1 reports the predictive performance of the evaluated graph-based models on the held-out test node subset $\mathcal{T} \subseteq V$, using root mean squared error (RMSE) and mean absolute error (MAE). Reported values correspond to the mean and standard deviation over multiple random seeds, with model selection performed using validation RMSE. For centralized training, validation performance is tracked per epoch, while for federated training it is tracked per communication round.

Table 1. Predictive performance comparison across centralized and federated baselines.

Model	RMSE (mean $\pm$ std)	MAE (mean $\pm$ std)
Centralized GAT	0.2123 $\pm$ 0.0805	0.1771 $\pm$ 0.0647
Federated GAT (FedAvg)	0.2836 $\pm$ 0.0892	0.2401 $\pm$ 0.0806
Federated Sustainability-Aware GAT ($\lambda = 3.0$)	0.2750 $\pm$ 0.0965	0.2313 $\pm$ 0.0810

The centralized GAT achieves the lowest prediction error across both metrics (RMSE = 0.2123 $\pm$ 0.0805, MAE = 0.1771 $\pm$ 0.0647), reflecting its unrestricted access to the full supply chain graph. This setting allows the model to exploit complete neighborhood structure and edge attributes across all tiers, which is particularly advantageous given the data-generating process. The synthetic node-level stress signal depends not only on node-local attributes but also on aggregated upstream quantities, such as total inbound flow and transport emissions, which are naturally captured through graph-based message passing.

The predictive advantage of centralized graph learning can be attributed to two complementary architectural properties. First, relational aggregation over incoming neighborhoods enables the model to smooth node representations by incorporating information from upstream entities $\mathcal{N}_{\text{in}}(v)$, which mitigates the effect of label noise by exploiting correlated structure across connected nodes. Second, edge-aware attention weighting allows the model to differentially weight upstream messages based on both node embeddings and normalized edge attributes (flow and transport emissions). This mechanism enables the model to identify and emphasize the most influential upstream connections for downstream stress propagation.

Both federated graph attention models exhibit higher prediction error than the centralized baseline, highlighting the intrinsic challenges of decentralized graph learning under strict locality constraints. The standard Federated GAT (FedAvg) attains $\text{RMSE} = 0.2836 \pm 0.0892$ and $\text{MAE} = 0.2401 \pm 0.0806$. In this setting, each client trains exclusively on its induced subgraph

$$G_k = (V_k, E_k),$$

and cross-client edges are unavailable during message passing. As a result, many true global dependencies are absent from local neighborhoods, and the global model obtained via weighted FedAvg can only reconcile these fragmented views indirectly through parameter averaging.

The Federated Sustainability-Aware GAT achieves slightly improved predictive performance relative to the standard federated baseline ($\text{RMSE} = 0.2750 \pm 0.0965$, $\text{MAE} = 0.2313 \pm 0.0810$). In this model, sustainability is not enforced through an auxiliary loss term but is instead embedded directly into the attention mechanism via an emission-aware bias applied after softmax normalization Equations (15) and (16):

With $\lambda = 3$, high-emission transport edges are systematically down-weighted during message passing. While this mechanism imposes an explicit inductive bias, it does not merely degrade predictive performance. Instead, the sustainability-aware formulation partially compensates for the loss of global visibility in federated training by guiding attention toward lower-emission, and often more structurally consistent, pathways within each client subgraph.

From an applied perspective, these results highlight a three-way trade-off. Centralized graph learning represents a best-case scenario for predictive accuracy but is often infeasible in practice due to data-sharing constraints. Standard federated learning preserves data locality but incurs a clear performance penalty when global relational dependencies are critical. The sustainability-aware federated GAT occupies an intermediate position: it remains fully decentralized while explicitly embedding environmental preferences into the inference mechanism, achieving competitive predictive performance alongside improved sustainability alignment. The resulting error should therefore be interpreted not as a limitation, but as the quantified cost of enforcing environmentally informed inductive biases within federated graph learning.

4.2. Training Dynamics and Convergence

Figure 3 presents the validation RMSE trajectories of the centralized and federated graph attention models, averaged over multiple random seeds, with shaded regions indicating one standard deviation. To ensure comparability across training paradigms, all curves are reported over a unified horizontal axis of 40 evaluation steps. For federated models, each step corresponds to one communication round, while for centralized training the validation trajectory is resampled from the full 250-epoch optimization history. The horizontal axis therefore represents aligned reporting checkpoints rather than equivalent optimization time.

The centralized GAT exhibits smooth and stable convergence, with a consistent reduction in validation RMSE and comparatively narrow variability across seeds. This behavior reflects the advantages of centralized graph learning, where full access to the global supply chain topology enables uninterrupted message passing across all tiers. At each optimization step, gradients are computed using complete neighborhood information and edge attributes, yielding low-variance updates and reliable convergence.

Both federated graph attention models converge more slowly and display wider uncertainty bands. This behavior is characteristic of decentralized graph learning under non-identically distributed (non-IID) subgraph partitions. Each client optimizes a local objective defined on its induced subgraph

$$G_k = (V_k, E_k),$$

where neighborhood structure, edge-attribute distributions, and structural roles within the supply chain differ across clients. These heterogeneous local optima are periodically reconciled through weighted parameter aggregation, which can introduce non-smooth global parameter updates and contributes to the increased variance observed in federated training. By construction, this model reshapes the message-passing process through an emission-aware attention bias that penalizes high-emission transport edges prior to renormalization. Consequently, information propagation is progressively concentrated on lower-emission pathways, leading to a more selective and structured flow of information during training. From an optimization standpoint, this mechanism introduces a controlled inductive bias that affects convergence speed but enhances consistency in how sustainability considerations are incorporated into the learned representations. In federated settings, where each client observes only a fragment of the global graph, unconstrained attention can overemphasize locally predictive but environmentally undesirable edges. The sustainability-aware formulation counteracts this tendency by steering optimization toward solutions that balance predictive performance with environmental alignment.




Figure 1. Validation RMSE trajectories for centralized and federated graph attention models (mean $\pm$ standard deviation).

The modest elevation in validation RMSE relative to the standard federated baseline should therefore be interpreted as the explicit cost of embedding sustainability constraints directly into the inference pathway, rather than as evidence of instability or inefficient training. Notably, the overlap of uncertainty bands in Figure 3b indicates that the sustainability-aware model remains stable across random initializations and does not exhibit pathological convergence behavior.

Overall, Figure 3b demonstrates that sustainability-aware federated graph learning converges reliably under strict locality constraints, while embedding environmental priorities directly into the message-passing mechanism. This positions the proposed model as a principled extension of federated graph learning for sustainability-critical supply chain applications, where predictive accuracy must be balanced against environmentally informed decision-making.

4.3. Sustainability-Aware Attention Patterns

A key objective of the proposed sustainability-aware federated GAT is to reduce reliance on carbon-intensive transport links during message passing. This behavior can be directly examined by

analyzing learned attention coefficients, extracted during evaluation with. For each directed edge ($u \rightarrow v$), attention weights are averaged across heads to obtain a single scalar α_{uv} .

Figure 4 shows the relationship between normalized transport emissions and mean attention weights for the standard Federated GAT ($\lambda = 0$) and the sustainability-aware variant ($\lambda = 3$). In the unconstrained model, attention weights are widely dispersed across the full emissions range, with no systematic dependence on emissions. Several high-emission edges receive moderate or high attention, reflecting their predictive utility under the synthetic stress generation mechanism, which explicitly incorporates inbound emissions.




Figure 2. Mean attention weight versus normalized transport emissions for standard and sustainability-aware federated GAT models.

In contrast, the sustainability-aware model exhibits a clear shift in attention allocation. High-emission edges are systematically down-weighted, while low- and medium-emission edges retain greater attention mass. This pattern is induced by the multiplicative emissions penalty applied within the attention mechanism, followed by neighborhood renormalization. The resulting behavior reflects a soft inductive bias: carbon-intensive links are discouraged but not entirely suppressed, allowing predictive relevance to partially offset the sustainability penalty when necessary.

From an interpretability standpoint, the learned attention distributions provide graph-native explanations of information flow. The observed reduction in attention assigned to high-emission edges confirms that sustainability considerations are embedded directly into the inference pathway, yielding representations that balance predictive performance with environmentally informed decision-making.

4.4. Ablation Study: Impact of Sustainability Weight (λ)

To systematically assess how the strength of sustainability bias influences the behavior of the proposed attention mechanism, we conduct a λ -ablation study over a predefined and fixed grid

$$\lambda \in \{0, 0.25, 0.5, 1, 2, 3, 5, 8\}, \quad (21)$$

with all configurations evaluated across multiple random seeds. The λ values are specified *a priori* and are not tuned post hoc, ensuring a transparent and reproducible evaluation protocol. While predictive performance is summarized separately, this subsection focuses specifically on how sustainability alignment evolves as a function of λ , isolating the effect of the attention penalty on message-passing behavior.

Figure 5 reports two complementary sustainability alignment diagnostics as functions of λ , shown as mean $\pm$ standard deviation over seeds:

- (i) the Spearman rank correlation between normalized transport emissions and learned attention weights, and
- (ii) the fraction of incoming attention mass allocated to high-emission edges (top emission quantile).




Figure 5. Effect of λ on Sustainability Alignment.

Both metrics exhibit smooth and monotonic trends as λ increases. The Spearman correlation becomes progressively more negative, indicating that attention weights are increasingly de-emphasized on high-emission edges. In parallel, the share of attention mass assigned to carbon-intensive links decreases steadily. These consistent patterns confirm that the proposed exponential attention penalty induces a systematic and stable reorientation of message passing away from high-emission transport links, rather than producing noisy or irregular effects.

Notably, the variability across random seeds decreases at higher λ values, suggesting that stronger sustainability bias leads to more reproducible alignment behavior under federated training conditions. This indicates that the mechanism functions as a stable inductive bias, rather than introducing additional instability into the learning process.

Figure 6 summarizes the ablation results by plotting predictive error against sustainability alignment, yielding a trade-off frontier across λ values. Each point corresponds to a distinct operating regime of the sustainability-aware federated GAT, with sustainability alignment measured as the negative Spearman correlation between emissions and attention (higher values indicate stronger de-emphasis of high-emission edges).




Figure 6. Accuracy–Sustainability Trade-Off.

The resulting frontier reveals a clear and structured trade-off. As λ increases, sustainability alignment improves monotonically, while predictive error increases gradually. Importantly, the degradation in accuracy is smooth rather than abrupt, indicating that the attention penalty suppresses high-emission links in a soft and continuous manner rather than eliminating large portions of the graph structure.

Intermediate λ values (approximately $\lambda \approx 2\text{--}3$) occupy favorable operating regimes near the knee of the trade-off curve. In this range, substantial reductions in emission–attention coupling are achieved while predictive degradation remains limited, representing practical compromises between environmental alignment and predictive fidelity.

Overall, this analysis reframes λ not as a hyperparameter to be optimized solely for accuracy, but as a policy-relevant control variable. Adjusting λ directly governs how strongly sustainability considerations shape the model’s inference pathway, enabling controlled, interpretable, and reproducible trade-offs in sustainability-critical federated supply chain settings.

5. Discussion

This study set out to address a critical gap at the intersection of supply chain analytics, federated learning, and sustainability-aware artificial intelligence. The results demonstrate that federated graph attention models can effectively learn over decentralized and partially observable supply chain networks, while explicitly incorporating environmental preferences into the inference process.

From a methodological perspective, the performance gap observed between centralized and federated GAT models confirms earlier findings that global graph visibility provides a structural advantage for relational learning. In the centralized setting, complete access to upstream and downstream dependencies enables efficient aggregation of operational and environmental signals, leading to lower prediction error. However, such centralized learning remains impractical in real-world supply chains due to confidentiality, competition, and regulatory constraints. The federated baselines therefore represent a more realistic deployment scenario, albeit at the cost of reduced structural observability.

Within this constrained setting, the sustainability-aware federated GAT demonstrates two key properties. First, it achieves predictive performance comparable to the standard federated GAT, despite embedding an explicit inductive bias against carbon-intensive transport links. This indicates

that sustainability alignment can be incorporated into decentralized graph learning without destabilizing training or causing disproportionate accuracy losses. Second, attention analysis confirms that the proposed mechanism systematically reshapes information flow, reducing the influence of high-emission edges during message passing. Unlike post hoc sustainability evaluation, this approach embeds environmental considerations directly into the model's internal reasoning, enabling inherently sustainability-aligned predictions.

The ablation study further reframes sustainability weighting not as a purely technical hyperparameter, but as a governance-relevant control variable. Adjusting the sustainability weight enables decision-makers to explicitly navigate the trade-off between operational accuracy and environmental alignment, producing a transparent accuracy–sustainability frontier. This is particularly relevant for regulatory and corporate contexts where sustainability targets must be balanced against service-level requirements.

Several limitations should be acknowledged. The experimental evaluation relies on a synthetic supply chain, which, while structurally realistic, cannot capture all complexities of real-world industrial networks. Additionally, the current framework focuses on node-level regression with static graph topology, whereas real supply chains exhibit temporal dynamics, evolving partnerships, and multi-modal data streams. Communication efficiency, security mechanisms, and adversarial robustness are also not explicitly addressed in this study and remain important considerations for large-scale deployment.

Overall, the findings suggest that sustainability-aware federated graph learning is both technically feasible and practically meaningful. By integrating environmental objectives directly into decentralized relational learning, the proposed framework advances the state of the art beyond accuracy-centric and centrally trained supply chain models.

6. Conclusions

This paper introduced a sustainability-aware federated graph attention framework for supply chain process modelling under strict data-governance constraints. By combining Graph Attention Networks with federated learning, the proposed approach enables collaborative learning across decentralized supply chain participants without sharing raw operational data. Sustainability considerations are embedded directly into the attention mechanism, guiding message passing away from carbon-intensive transport links.

Experimental results on a multi-tier synthetic supply chain demonstrate that while centralized graph learning remains an upper bound in predictive accuracy, the proposed sustainability-aware federated GAT achieves competitive performance relative to standard federated baselines. Importantly, it delivers measurable sustainability alignment through interpretable attention patterns, rather than treating environmental metrics as external evaluation criteria. The ablation analysis further shows that sustainability–accuracy trade-offs can be controlled smoothly and transparently through a single model parameter.

These findings contribute to both theory and practice by showing how graph-based, privacy-preserving, and sustainability-aware AI can be jointly realized in supply chain settings. Future work will focus on extending the framework to dynamic and temporal graphs, integrating multiple environmental indicators, and validating the approach on real industrial datasets. More broadly, this work positions federated graph attention models as a promising foundation for trustworthy and sustainability-aligned AI in multi-enterprise supply chains

Author Contributions: V.A.: conceptualization, data curation, formal analysis, investigation, methodology, resources, software, validation, visualization, writing original draft, and writing review and editing. M.D.: conceptualization, investigation, methodology, project administration, resources, supervision, validation, writing original draft, and writing review and editing. P.T.: supervision, validation, and writing review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Data Availability Statement: No new data were created.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
CPN	Colored Petri Net(s)
CPU	Central Processing Unit
CO ₂	Carbon Dioxide
FedAvg	Federated Averaging
FedGNN	Federated Graph Neural Network
FL	Federated Learning
GAT	Graph Attention Network
GCN	Graph Convolutional Network
GDPR	General Data Protection Regulation
GNN	Graph Neural Network
GPU	Graphics Processing Unit
ICT	Information and Communication Technology
IID	Independent and Identically Distributed
IoT	Internet of Things
KPI	Key Performance Indicator
LR	Learning Rate
MAE	Mean Absolute Error
MSE	Mean Squared Error
Non-IID	Non-Independent and Identically Distributed
PyTorch	Python-based deep learning framework (PyTorch)
ReLU	Rectified Linear Unit
RMSE	Root Mean Squared Error
VFGNN	Vertically Federated Graph Neural Network
WD	Weight Decay

References

1. Drakaki, M.; Tzionas, P. Modelling and Performance Evaluation of an Agent-Based Warehouse Dynamic Resource Allocation Using Colored Petri Nets. *Int. J. Comput. Integr. Manuf.* 2016, 29 (7), 736–753. <https://doi.org/10.1080/0951192X.2015.1130239>.
2. Drakaki M, Tzionas P. A colored petri net-based modelling method for supply chain inventory management. *SIMULATION*. 2021;98(3):257-271. doi:10.1177/00375497211038755
3. Drakaki, M.; Tzionas, P. Investigating the Impact of Inventory Inaccuracy on the Bullwhip Effect in RFID-Enabled Supply Chains Using Colored Petri Nets. *J. Model. Manag.* 2019, 14 (2), 360–384. <https://doi.org/10.1108/JM2-08-2017-0081>.
4. Drakaki, M.; Tzionas, P. Manufacturing Scheduling Using Colored Petri Nets and Reinforcement Learning. *Appl. Sci.* 2017, 7 (2). <https://doi.org/10.3390/app7020136>.
5. Drakaki, M.; Gören, H.; Tzionas, P. A Multi-Agent Based Decision Framework for Sustainable Supplier Selection, Order Allocation and Routing Problem: In *Proceedings of the 5th International Conference on Vehicle Technology and Intelligent Transport Systems; SCITEPRESS - Science and Technology Publications: Crete, Greece, 2019; pp 621–628.* <https://doi.org/10.5220/0007833306210628>.
6. Drakaki, M.; Gören, H. G.; Tzionas, P. Supplier Selection Problem in Fuzzy Environment Considering Risk Factors. In *2019 12th International Conference on Developments in eSystems Engineering (DeSE); 2019; pp 784–788.* <https://doi.org/10.1109/DeSE.2019.00146>.

7. Drakaki, M.; Karnavas, Y. L.; Karlis, A. D.; Chasiotis, I. D.; Tzionas, P. Study on Fault Diagnosis of Broken Rotor Bars in Squirrel Cage Induction Motors: A Multi-Agent System Approach Using Intelligent Classifiers. *IET Electr. Power Appl.* 2020, 14 (2), 245–255. <https://doi.org/10.1049/iet-epa.2019.0619>.
8. Drakaki, M.; Karnavas, Y. L.; Karlis, A. D.; Chasiotis, I. D.; Tzionas, P. Study on Fault Diagnosis of Broken Rotor Bars in Squirrel Cage Induction Motors: A Multi-Agent System Approach Using Intelligent Classifiers. *IET Electr. Power Appl.* 2020, 14 (2), 245–255. <https://doi.org/10.1049/iet-epa.2019.0619>.
9. Neural Networks in Supply Chain Analytics and Optimization: Concepts, Perspectives, Dataset and Benchmarks, 2024. <https://doi.org/10.5281/ZENODO.13652826>.
10. Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; Bengio, Y. Graph Attention Networks. *arXiv* February 4, 2018. <https://doi.org/10.48550/arXiv.1710.10903>.
11. S.Kayalvizhi. Graph Neural Network-Based Predictive Modelling for Enhanced Supply Chain Resilience against Multi-Modal Disruptions. *J. Inf. Syst. Eng. Manag.* 2025, 10 (33s), 628–642. <https://doi.org/10.52783/jisem.v10i33s.5571>.
12. Tang, X.; Hu, F.; Tang, X.; Hu, F. A Privacy-Enhancing Mechanism for Federated Graph Neural Networks. *Symmetry* 2025, 17 (4). <https://doi.org/10.3390/sym17040565>.
13. McMahan, H. B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B. A. y. Communication-Efficient Learning of Deep Networks from Decentralized Data. *arXiv* January 26, 2023. <https://doi.org/10.48550/arXiv.1602.05629>.
14. (Islam, F.; Raihan, A. S.; Ahmed, I. Applications of Federated Learning in Manufacturing: Identifying the Challenges and Exploring the Future Directions with Industry 4.0 and 5.0 Visions. *arXiv* May 16, 2023. <https://doi.org/10.48550/arXiv.2302.13514>.
15. Frikha, M. A.; Mrad, M.; Frikha, M. A.; Mrad, M. AI-Driven Supply Chain Decarbonization: Strategies for Sustainable Carbon Reduction. *Sustainability* 2025, 17 (21). <https://doi.org/10.3390/su17219642>.
16. Sievers, J.; Henrich, P.; Beichter, M.; Mikut, R.; Hagenmeyer, V.; Blank, T.; Simon, F. Federated Reinforcement Learning for Sustainable and Cost-Efficient Energy Management. *Energy AI* 2025, 21, 100521. <https://doi.org/10.1016/j.egyai.2025.100521>.
17. Zheng, G.; Brintrup, A. A Machine Learning Approach for Enhancing Supply Chain Visibility with Graph-Based Learning. *Supply Chain Anal.* 2025, 11, 100135. <https://doi.org/10.1016/j.sca.2025.100135>.
18. Abushaega, M. M.; Moshebah, O. Y.; Hamzi, A.; Alghamdi, S. Y. Multi-Objective Sustainability Optimization in Modern Supply Chain Networks: A Hybrid Approach with Federated Learning and Graph Neural Networks. *Alex. Eng. J.* 2025, 115, 585–602. <https://doi.org/10.1016/j.aej.2024.12.048>.
19. Gupta, I.; Martinez, A.; Correa, S.; Wicaksono, H. A Comparative Assessment of Causal Machine Learning and Traditional Methods for Enhancing Supply Chain Resiliency and Efficiency in the Automotive Industry. *Supply Chain Anal.* 2025, 10, 100116. <https://doi.org/10.1016/j.sca.2025.100116>.
20. Lee, D.; Go, J.; Noh, T.; Song, S. Multi-Feature Representation-Based Graph Attention Networks for Predicting Potential Supply Relationships in a Large-Scale Supply Chain Network. *Expert Syst. Appl.* 2025, 292, 128593. <https://doi.org/10.1016/j.eswa.2025.128593>.
21. Park, S.; Lee, H. Predictive Supply Chain Disruption Control Framework Using Casual Network-Based Multi-Stream Deep Learning. *Comput. Ind. Eng.* 2025, 207, 111312. <https://doi.org/10.1016/j.cie.2025.111312>.
22. Yan Hong Lim, K.; Liu, Y.; Chen, C.-H.; Gu, X. Manufacturing Resilience through Disruption Mitigation Using Attention-Based Consistently-Attributed Graph Embedded Decision Support System. *Comput. Ind. Eng.* 2024, 197, 110494. <https://doi.org/10.1016/j.cie.2024.110494>.
23. Kosasih, E. E.; Brintrup, A. A Machine Learning Approach for Predicting Hidden Links in Supply Chain with Graph Neural Networks. *Int. J. Prod. Res.* 2022, 60 (17), 5380–5393. <https://doi.org/10.1080/00207543.2021.1956697>.
24. Cai, C.; Pan, L.; Li, X.; Luo, S.; Wu, Z. A Risk Identification Model for ICT Supply Chain Based on Network Embedding and Text Encoding. *Expert Syst. Appl.* 2023, 228, 120459. <https://doi.org/10.1016/j.eswa.2023.120459>.
25. Ahmed, I.; Ahmad, M.; Jeon, G. Federated Learning in Convergence ICT: A Systematic Review on Recent Advancements, Challenges, and Future Directions. *Comput. Mater. Contin.* 2025, 85 (3), 4237–4273. <https://doi.org/10.32604/cmc.2025.068319>.

26. Takele, A. K.; Villányi, B.; Takele, A. K.; Villányi, B. Resource-Efficient Clustered Federated Learning Framework for Industry 4.0 Edge Devices. *AI* 2025, 6 (2). <https://doi.org/10.3390/ai6020030>.

27. Leng, J.; Li, R.; Xie, J.; Zhou, X.; Li, X.; Liu, Q.; Chen, X.; Shen, W.; Wang, L. Federated Learning-Empowered Smart Manufacturing and Product Lifecycle Management: A Review. *Adv. Eng. Inform.* 2025, 65, 103179. <https://doi.org/10.1016/j.aei.2025.103179>.

28. Makanda, I. L. D.; Jiang, P.; Yang, M. Personalized Federated Unsupervised Learning for Nozzle Condition Monitoring Using Vibration Sensors in Additive Manufacturing. *Robot Comput-Integr Manuf* 2025, 93 (C). <https://doi.org/10.1016/j.rcim.2024.102940>.

29. Li, D.; Liu, S.; Wang, B.; Yu, C.; Zheng, P.; Li, W. Trustworthy AI for Human-Centric Smart Manufacturing: A Survey. *J. Manuf. Syst.* 2025, 78, 308–327. <https://doi.org/10.1016/j.jmsy.2024.11.020>.

30. Federated Learning-Based Framework: A New Paradigm Proposed for Supply Chain Risk Management | MDPI. <https://www.mdpi.com/2673-4591/97/1/5> (accessed 2026-01-11).

31. Liang, X. Cross-Border Logistics Risk Warning System Based on Federated Learning. *Sci. Rep.* 2025, 15 (1), 39131. <https://doi.org/10.1038/s41598-025-25507-1>.

32. Gavai, A. K.; Bouzembrak, Y.; Xhani, D.; Sedrakyan, G.; Meuwissen, M. P. M.; Souza, R. G. S.; Marvin, H. J. P.; Hillegersberg, J. van. Agricultural Data Privacy: Emerging Platforms & Strategies. *Food Humanity* 2025, 4, 100542. <https://doi.org/10.1016/j.foohum.2025.100542>.

33. Fendor, Z.; Velden, B. H. M. van der; Wang, X.; Carnoli, A. J.; Mutlu, O.; Hürriyetoğlu, A. Federated Learning in Food Research. *J. Agric. Food Res.* 2025, 23, 102238. <https://doi.org/10.1016/j.jafr.2025.102238>.

34. Tang, X.; Wang, Y.; Liu, X.; Yuan, X.; Fan, C.; Hu, Y.; Miao, Q. Federated Graph Neural Network for Privacy-Preserved Supply Chain Data Sharing. *Appl Soft Comput* 2025, 168 (C). <https://doi.org/10.1016/j.asoc.2024.112475>.

35. Benzidia, S.; Makaoui, N.; Bentahar, O. The Impact of Big Data Analytics and Artificial Intelligence on Green Supply Chain Process Integration and Hospital Environmental Performance. *Technol. Forecast. Soc. Change* 2021, 165, 120557. <https://doi.org/10.1016/j.techfore.2020.120557>.

36. Yadav, A.; Garg, R. K.; Sachdeva, A. Artificial Intelligence Applications for Information Management in Sustainable Supply Chain Management: A Systematic Review and Future Research Agenda. *Int. J. Inf. Manag. Data Insights* 2024, 4 (2), 100292. <https://doi.org/10.1016/j.jjime.2024.100292>.

37. Alharithi, F. S.; Alzahrani, A. A. Enhancing Environmental Sustainability with Federated LSTM Models for AI-Driven Optimization. *Alex. Eng. J.* 2024, 108, 640–653. <https://doi.org/10.1016/j.aej.2024.09.058>.

38. Alexiadis, V.; Drakaki, M.; Tzionas, P.; Alexiadis, V.; Drakaki, M.; Tzionas, P. LSTM-Based Electricity Demand Forecasting in Smart and Sustainable Hospitality Buildings. *Electronics* 2025, 14 (22). <https://doi.org/10.3390/electronics14224456>.

39. Borraccia, S.; Masciari, E.; Napolitano, E. V. Green Metrics for AI: A Hybrid Strategy for Environmental Impact Assessment. *Array* 2025, 28, 100528. <https://doi.org/10.1016/j.array.2025.100528>.

40. Valencia-Arias, A.; Vásquez Coronado, M. H.; Clavo Medina, M. L.; Morales Garcia, M. J.; Requejo Arias, A. R.; Saavedra Tirado, J. I. Research Trends in the Application of Machine Learning in Sustainability Practices Based on a Bibliometric Analysis. *Sustain. Futur.* 2025, 10, 100987. <https://doi.org/10.1016/j.sftr.2025.100987>.

41. Nti, E. K.; Cobbina, S. J.; Attafuah, E. E.; Opoku, E.; Gyan, M. A. Environmental Sustainability Technologies in Biodiversity, Energy, Transportation and Water Management Using Artificial Intelligence: A Systematic Review. *Sustain. Futur.* 2022, 4, 100068. <https://doi.org/10.1016/j.sftr.2022.100068>.

42. Zechiel, F.; Blaurock, M.; Weber, E.; Büttgen, M.; Coussement, K. How Tech Companies Advance Sustainability through Artificial Intelligence: Developing and Evaluating an AI x Sustainability Strategy Framework. *Ind. Mark. Manag.* 2024, 119, 75–89. <https://doi.org/10.1016/j.indmarman.2024.03.010>.

43. AI cuts supply chain emissions by over 1,000 tons annually. <https://www.devdiscourse.com/article/technology/3678443-ai-cuts-supply-chain-emissions-by-over-1000-tons-annually> (accessed 2026-01-17).
44. Slimani, S.; Omri, A.; Ben Jabeur, S. When and How Does Artificial Intelligence Impact Environmental Performance? *Energy Econ.* 2025, 148, 108643. <https://doi.org/10.1016/j.eneco.2025.108643>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.